



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/643,588	08/18/2003	Kitrick Sheets	1376.720US1	4010
21186	7590	04/08/2008		
SCHWEGMAN, LUNDBERG & WOESSNER, P.A. P.O. BOX 2938 MINNEAPOLIS, MN 55402			EXAMINER	
			TSAL, SHIENG JEN	
ART UNIT		PAPER NUMBER		
2186				
MAIL DATE		DELIVERY MODE		
04/08/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 10/643,588	Applicant(s) SHEETS, KITRICK
	Examiner SHENG-JEN TSAI	Art Unit 2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If no period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED. (35 U.S.C. § 133).

Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 17 March 2008.

2a) This action is FINAL. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1,4-6,9-11 and 14-18 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) Claim(s) _____ is/are allowed.

6) Claim(s) 1,4-6,9-11 and 14-18 is/are rejected.

7) Claim(s) _____ is/are objected to.

8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on 18 August 2003 is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a) All b) Some * c) None of:

1. Certified copies of the priority documents have been received.
2. Certified copies of the priority documents have been received in Application No. _____.
3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)

2) Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____

4) Interview Summary (PTO-413)
Paper No(s)/Mail Date _____

5) Notice of Informal Patent Application

6) Other: _____

DETAILED ACTION

1. This Office Action is taken in response to Applicant's Request for Continued Examination (RCE) filed on March 17, 2008 regarding application 10/643,588 filed on August 18, 2003.

2. Claims 2-3, 7-8 and 12-13 have been cancelled previously.

Claims 1, 4-6, 9-11 and 14-18 are pending under consideration.

3. **Response to Remarks**

Applicant's amendments and remarks have been fully and carefully considered.

In response to the amendments, a new ground of claim analysis based on references Scott (US Patent Application Publication 2004/0044872) and Schimmel (US 6,105,113) has been made. Refer to the corresponding sections of claim analysis for details.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 1, 4-6, 9-11 and 14-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Scott (US Patent Application Publication 2004/0044872), and in view of Schimmel (US 6,105,113).

As to claim 1, Scott discloses a method for translating a virtual memory address into a physical memory address in a multi-node system [Remote

Translation Mechanism for a Multi-Node System (title); A method for remotely translating a virtual memory address into a physical memory address in a multi-node system (claim 1); figure 1 shows a multi-node system; a remote translation mechanism for a multi-node system (abstract)], **the method comprising:**

Initializing in a generally accessible memory [figure 5 shows that the RTT comprising a 64K entries memory; figures 6A~6D show the memory organization of the RTT; translating the virtual memory address on the remote node into a physical memory address using a remote-translation table (RTT). The RTT contains translation information for an entire virtual memory address space associated with the remote node (abstract); thus RTT is a generally accessible memory because translations are written into it and read from it to facilitate translating the virtual memory address into a physical memory address; in order for the remote translation mechanism disclosed by Scott to work and function properly, it is inherent that the RTT at all the nodes be initialized and synchronized first before any reference to a memory location resides at a remote node can be served. Without the initialization and synchronization, the RTT may not have the correct information to reach the correct memory location (this also applies to the system disclosed by Schimmel)] **an emulated remote translation table (ERTT) segment** [figures 6A~6D show the memory organization of the RTT including the segments (6D); translating the virtual memory address on the remote node into a physical memory address using a remote-translation table (RTT). The RTT contains translation information for an entire virtual memory address space associated with the remote node (abstract)];

providing the virtual memory address at a source node [providing the virtual memory address at a source node (claim 1)];

determining that a translation for the virtual memory address does not exist

[determining that the virtual memory address is to be sent to a remote node (claim 1);

wherein the determining that the virtual memory address is to be sent to the remote

node includes determining that a virtual node field does not match the source node

(claim 2); determining if the virtual node corresponds to the local node; translating the

virtual memory address into a local physical memory address on the local node, if the

virtual node corresponds to the local node; and if the virtual node corresponds instead

to a remote node, sending the virtual memory address to the remote node, and

translating the virtual memory address into a physical memory address on the remote

node (claim 14)];

determining a virtual node to query based on the virtual memory address

[wherein the determining that the virtual memory address is to be sent to the remote

node includes determining that a virtual node field does not match the source node

(claim 2); figure 4; paragraph 0031; in one implementation, a local node can identify

the virtual node by looking at the VNode field of the virtual address. Checkpoint 404

determines if the virtual node is the same as, or equal to, the local node. If so, flow

diagram 400 continues to block 406, wherein the virtual address is translated into a

physical address locally using a Translation Look-Aside Buffer (TLB). The local node is

then able to address local physical memory space. If the virtual node is not the same

as the local node, then flow diagram 400 continues to block 408, wherein the virtual

address is translated into a physical address remotely (on a remote node) using a Remote-Translation Table (RTT). In this fashion, the local node is effectively able to address remote memory space of the remote node (paragraph 0031)];

accessing an ERTT header to obtain a mapping the virtual node to a physical node [wherein the using of the virtual node identifier includes combining the virtual node identifier with a base node identifier to determine the physical node identifier of the remote node (claim 8); figure 5 illustrates how this mapping is accomplished staring from the Vnode (virtual node) specified by the source node and ending with the corresponding physical address at a local (remote) node];

querying the ERTT segment on the physical node for the translation for the virtual memory address [wherein the combining includes adding the virtual node identifier to the base node identifier to determine the physical node identifier of the remote node, wherein the base node identifier is maintained by, and unique to, the source node (claim 9); wherein the using of the virtual node identifier includes using the virtual node identifier as an index into a look-up table to determine the physical node identifier of the remote node (claim 10); figure 5 illustrates how this mapping is accomplished staring from the Vnode (virtual node) specified by the source node and ending with the corresponding physical address at a local (remote) node; the corresponding ERTT is the Global Address Space identifier (GASID) and the Remote Translation Table (RTT) as described in paragraph 0019]; **and**
loading the translation into a translation lookaside buffer (TLB) on the source node [taught by Schimmel, see below].

Regarding claim 1, Scott does not teach **loading the translation into a translation lookaside buffer (TLB) on the source node.**

However, Schimmel teaches this limitation in the invention "System and Method for Maintaining Translation Look-Aside Table (TLB) Consistency," a system and method of maintaining consistency of a TLB which translates virtual memory addresses into the corresponding physical addresses [abstract].

Particularly, figure 9 of Schimmel teaches that, if there is a TLB miss (NO at step 916), the corresponding translation will be obtained (steps 922, 924 and 926), sent to the source processor/node (step 928), and loaded into the TLB of the source processor/node (step 930).

Schimmel also teaches the motivation of loading translation into the TLB of the source processor/node is to reduce the latency of accessing the translation [col. 1, line 48 to col. 2, line 10]

Therefore, it would have been obvious for one of ordinary skills in the art at the time of Applicant's invention to load translation into the TLB of the source processor/node, as demonstrated by Schimmel, in order to reduce the latency of accessing the translation.

As to claim 4, Scott teaches that **mapping the virtual node to a physical node uses a mapping provided by an ERTT header located at a well-known location to one or more nodes used by an application** [wherein the method further comprises sending a global address space identifier (GASID) to the remote node along with the virtual memory address, and wherein the translating includes using the GASID to

translate the virtual memory address into a physical memory address (claim 6); abstract; figure 5 illustrates how this mapping is accomplished starting from the Vnode (virtual node) specified by the source node and ending with the corresponding physical address at a local (remote) node].

As to claim 5, Scott teaches that **the ERTT header is located on a predetermined virtual node** [determining that the virtual memory address is to be sent to a remote node (claim 1); wherein the determining that the virtual memory address is to be sent to the remote node includes determining that a virtual node field does not match the source node (claim 2); determining if the virtual node corresponds to the local node; translating the virtual memory address into a local physical memory address on the local node, if the virtual node corresponds to the local node; and if the virtual node corresponds instead to a remote node, sending the virtual memory address to the remote node, and translating the virtual memory address into a physical memory address on the remote node (claim 14); figure 5 shows where the ERTT is located with respect to the source node. Since none of the node has all the translation information needed by itself, the collection of all the memories appears to be a predetermined virtual node as far as each of the node is concerned].

As to claim 6, Scott discloses a computerized system for managing virtual address translations, the system comprising:
a plurality of nodes [figure 1 shows a multi-node system] **available for executing programs** [figure 4 shows the computer programs], **each of said nodes having a node memory** [figure 1]; and

an operating system [it is common knowledge that all computers and processors are equipped with an operating system; the OS must never create a partition (via the BaseNode and NodeLimit values) that exceeds the number of nodes in the machine (paragraph 0033)] **executable by a source node of the plurality of nodes, the operating system operable to:**

receiving a virtual memory address at a source node [providing the virtual memory address at a source node (claim 1)];

determining that a translation for the virtual memory address does not exist on the source node [determining that the virtual memory address is to be sent to a remote node (claim 1); wherein the determining that the virtual memory address is to be sent to the remote node includes determining that a virtual node field does not match the source node (claim 2); determining if the virtual node corresponds to the local node; translating the virtual memory address into a local physical memory address on the local node, if the virtual node corresponds to the local node; and if the virtual node corresponds instead to a remote node, sending the virtual memory address to the remote node, and translating the virtual memory address into a physical memory address on the remote node (claim 14)];

determining a virtual node to query based on the virtual memory address [wherein the determining that the virtual memory address is to be sent to the remote node includes determining that a virtual node field does not match the source node (claim 2); figure 4; paragraph 0031; in one implementation, a local node can identify the virtual node by looking at the VNode field of the virtual address. Checkpoint 404

determines if the virtual node is the same as, or equal to, the local node. If so, flow diagram 400 continues to block 406, wherein the virtual address is translated into a physical address locally using a Translation Look-Aside Buffer (TLB). The local node is then able to address local physical memory space. If the virtual node is not the same as the local node, then flow diagram 400 continues to block 408, wherein the virtual address is translated into a physical address remotely (on a remote node) using a Remote-Translation Table (RTT). In this fashion, the local node is effectively able to address remote memory space of the remote node (paragraph 0031)];

accessing an ERTT header to obtain a mapping the virtual node to a physical node [wherein the using of the virtual node identifier includes combining the virtual node identifier with a base node identifier to determine the physical node identifier of the remote node (claim 8); figure 5 illustrates how this mapping is accomplished starting from the Vnode (virtual node) specified by the source node and ending with the corresponding physical address at a local (remote) node];

querying the ERTT segment in the generally accessible memory [figure 5 shows that the RTT comprising a 64K entries memory; figures 6A~6D show the memory organization of the RTT; translating the virtual memory address on the remote node into a physical memory address using a remote-translation table (RTT). The RTT contains translation information for an entire virtual memory address space associated with the remote node (abstract); thus RTT is a generally accessible memory because translations are written into it and read from it to facilitate translating the virtual memory address into a physical memory address] **on the physical node for the translation**

for the virtual memory address [wherein the combining includes adding the virtual node identifier to the base node identifier to determine the physical node identifier of the remote node, wherein the base node identifier is maintained by, and unique to, the source node (claim 9); wherein the using of the virtual node identifier includes using the virtual node identifier as an index into a look-up table to determine the physical node identifier of the remote node (claim 10); figure 5 illustrates how this mapping is accomplished starting from the Vnode (virtual node) specified by the source node and ending with the corresponding physical address at a local (remote) node; the corresponding ERTT is the Global Address Space identifier (GASID) and the Remote Translation Table (RTT) as described in paragraph 0019]; **and**
loading the translation into a translation lookaside buffer (TLB) on the source node [taught by Schimmel, see below].

Regarding claim 6, Scott does not teach **loading the translation into a translation lookaside buffer (TLB) on the source node**.

However, Schimmel teaches this limitation in the invention "System and Method for Maintaining Translation Look-Aside Table (TLB) Consistency," a system and method of maintaining consistency of a TLB which translates virtual memory addresses into the corresponding physical addresses [abstract].

Particularly, figure 9 of Schimmel teaches that, if there is a TLB miss (NO at step 916), the corresponding translation will be obtained (steps 922, 924 and 926), sent to the source processor/node (step 928), and loaded into the TLB of the source processor/node (step 930).

Schimmel also teaches the motivation of loading translation into the TLB of the source processor/node is to reduce the latency of accessing the translation [col. 1, line 48 to col. 2, line 10]

Therefore, it would have been obvious for one of ordinary skills in the art at the time of Applicant's invention to load translation into the TLB of the source processor/node, as demonstrated by Schimmel, in order to reduce the latency of accessing the translation.

As to claim 9, it recites substantially the same limitations as in claim 4, and is rejected for the same reasons set forth in the analysis of claim 4. Refer to "As to claim 4" presented earlier in this Office Action for details.

As to claim 10, it recites substantially the same limitations as in claim 5, and is rejected for the same reasons set forth in the analysis of claim 5. Refer to "As to claim 5" presented earlier in this Office Action for details.

As to claim 11, it recites substantially the same limitations as in claim 1, and is rejected for the same reasons set forth in the analysis of claim 1. Refer to "As to claim 1" presented earlier in this Office Action for details.

As to claim 14, it recites substantially the same limitations as in claim 4, and is rejected for the same reasons set forth in the analysis of claim 4. Refer to "As to claim 4" presented earlier in this Office Action for details.

As to claim 15, it recites substantially the same limitations as in claim 5, and is rejected for the same reasons set forth in the analysis of claim 5. Refer to "As to claim 5" presented earlier in this Office Action for details.

As to claim 16, Scott teaches **the method of claim 1, further comprising replicating the ERTT header on a plurality of physical nodes** [as shown in figures 1 and 6A, figure 1 shows a plurality of M chips (M0 through M15) and figure 6A shows that each M chip has a RTT (601)].

As to claim 17, Scott teaches **the system of claim 9, further comprising a plurality of replicated ERTT headers provided on a plurality of physical nodes** [as shown in figures 1 and 6A, figure 1 shows a plurality of M chips (M0 through M15) and figure 6A shows that each M chip has a RTT (601)].

As to claim 18, Scott teaches **the computer-readable medium of claim 14, wherein the method further comprises replicating the ERTT header on a plurality of physical nodes** [as shown in figures 1 and 6A, figure 1 shows a plurality of M chips (M0 through M15) and figure 6A shows that each M chip has a RTT (601)].

6. Related Prior Art

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Scott et al. (US 6,925,547), "Remote Address Translation in a Multiprocessor System."
- Deneau, (US 6,684,305), "Multiprocessor System Implementing Virtual Memory Using a Shared Memory, and a Page Replacement Method for Maintaining Paged memory Coherence."

- Frank et al., (US 6,490,671), "System for Efficiently Maintaining Translation Lookaside Buffer Consistency in a Multi-Threaded, Multi-Processor Virtual Memory System."
- Hansen, (US 6,101,590), "Virtual Memory System with Local and Global Virtual Address Translation."

Conclusion

7. Claims 1, 4-6, 9-11 and 14-18 are rejected as explained above.
8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Sheng-Jen Tsai/
Partial Signatory Examiner, Art Unit 2186

April 1, 2008

Application/Control Number: 10/643,588
Art Unit: 2186

Page 14